HHAI 2023: Augmenting Human Intellect
P. Lukowicz et al. (Eds.)

doi:10.3233/FAIA230129

Neurosymbolic Narrative Generation for Cultural Heritage

Cosimo PALMA
[a] *University of Naples "L'Orientale"*
[b] *University of Pisa*
ORCiD ID: Cosimo Palma https://orcid.org/0000-0002-8161-9782

Abstract. Aim of my research is to exploit Linguistic Linked Open Data (LLOD) as base for advanced Cultural Heritage (CH) fruition by means of Automatic Story Generation (ASG) . Following the rationale that discovering and reviving already existing (yet latent) narratives is worthier than automatically generating them from anew in eliciting the user's interest, the *input-2-graph* and the *graph-2-sequence* ASG-pipeline phases, heavily relying on LLOD, will be given a deeper focus, whereby the final Natural Language Generation (NLG) module will be constrained by the entities and relations established in the Knowledge Graph (KG) generation modules (a configuration typical of the *neurosymbolic* approach). In order to enhance possibilities of implementation in real-life contexts, the elaborated pipeline will be modular, i.e. self-sufficient in its constituent parts. Beyond the countless possible application scenarios ranging from education to entertainment, this solution detangles the user from his role of mere consumer, and empowers him not only to control the creation process [3.1], but also to find already within it, and not necessarily in the final outcome, a valuable source for intellectual growth. This work intends the addressing of a specific societal need as an avalanche to simultaneously fill knowledge gaps identified in and among the related scientific domains.

Keywords. Neurosymbolic AI, Linguistic Linked Open Data, Cultural Heritage, Automatic Story Generation, Knowledge Graphs

1. Introduction

This thesis explores neurosymbolic approaches to ASG, leveraging Linked Open Data (LOD, in particular the linguistic ones) and LOD-aware Natural Language Generation (NLG) procedures. More precisely, the research will encompass the development of a (semi-)automated narrative generation pipeline for digital heritage[1] drawing from semantically enriched data [1,2,3]. Stories will be entailed from them building upon the relations among their respective metadata.

With "semi-automation" is meant that all intermediate steps leading to the final outcome will be modular, i.e. already designed to be self-sufficient, easily interpretable and ready to be used independently from one another (their eventual composition should occur manually) at least as an aid to story creation, following the principle that the direct

[1] A domain which besides data availability shows a particular need for content creation as well.

intervention of the user is desirable for his engagement[2].
The related work [2] boiled down from the problem statement lets emerge many areas for improvement, pointing to a main resulting intuition, which my research will contribute to test:

Interestingness and coherence may be strongly improved by deepening the focus on the input-2-graph *and the* graph-2-sequence *ASG-pipeline phases.*

At the same time this statement also technically implements the rationale that discovering and reviving already existing (yet latent) narratives is worthier than automatically generating them from anew in eliciting the user's interest. The concept of Linked Open Data relates immediately to that of *network* (or graph), made up of links and nodes.

CH is one of the domains which has experienced the most massive impact of digitalization. Countless other ways of igniting a playful synergy with the user is object of current research, and has already delivered encouraging results [5,6]; for instance, plenty of systems interacting with the user by natural language are flourishing, either in form of a smartphone application [7], or even as social robot [8].
The main factors determining this contamination are the huge amount of information resources, and the rising need for alternative methods to engage with the public, whose attention's quality has shifted considerably during the last decades [9]. The decrease in the focus span given to the exponential growth of entertainment stimuli, as well the gradual detachment of the audience from traditional fruition of cultural heritage, has highlighted the urge to counteract this tendency, by putting into play various strategies, ranging from the transmedial narratives [10] to gamification [11], from Virtual Reality to interactive story-telling [12,11,13]. For this reason this work aims at the generation of narratives. The *narrative*, here used as a synonym for *story*, differently from a relevant portion of related literature, is in the following intended in its original sense, i.e. as an imaginative account of events involving either real or fictitious characters, places and times, designed to interest, amuse and educate [3]. The term "narrative" has been preferred because even "story" is subject to similar ambiguities; moreover, it better conveys the idea of a presentation, whereby the objective substance is deployed in a particular fashion from a specific standpoint. The term "story", on the other hand, seems to hint more specifically to a plot, and to literature features such as the writing style.

Currently in the cultural heritage domain, the single institutions provide their open data to the national aggregator which can forward it to the European database for CH known as *Europeana*[4] [14,15].

Network science's tools and concepts can be effectively used for exploring and comparing semantic spaces of word embeddings and lexical databases as well. Specifically, semantic networks based on *word2vec*[5] representation of words have shown that although human built networks possess more intuitive global connectivity patterns, local

[2]This work may also seen as an attempt of automatizing the genesis of narratives as Benjamin's *The Arcades Project*: "Method of this project: *literary montage*. I needn't say anything, merely show. I shall purloin no valuables, appropriate no ingenious formulations. But the rags, the refuse – these I will not inventory but allow, in the only way possible, to come into their own: by making use of them." [4].

[3]Freely adapted from *https://www.dictionary.com/browse/narrative* .

[4]https://www.europeana.eu/it. The data used for this work will be harvested principally from Europeana as well.

[5]https://www.tensorflow.org/tutorials/text/word2vec.

characteristics (in particular, *dense clusters*) of the machine built networks provide much richer information on the contextual usage and perceived meanings of words [16].
New frameworks for automatic KG construction empowered by the neural Language Models' flexibility and scalability have been established, requiring as input only the minimal definition of relations, and hence resulting fit for extracting knowledge of rich new relations not available before (a task called *Knowledge Discovery*) [17]. This technique would empower us to populate the KG, whose sequentialization would constitute the skeleton on top of which the story is generated. Conversely, semantically enriched data can foster the shrinking of the amount of training data necessary to learn accurate models, thus bringing to Machine Learning the concrete chance to achieve Few-Shot-Learning in a variety of tasks, such as event extraction, link prediction and KG embedding [18]. Although we are operating in a scenario where data is semantically enriched, the probable related quality and quantity scarcity may result in the need to use these type of technologies. For instance, The BiographyNet project shows how existing Linked Data vocabularies can be re-used for tasks as object modelling, resulting in a better compatibility of the data with other sources, especially datasets from *Europeana*. Furthermore, the use of Linked Data allowed to gradually expand the data corpus, which originally consisted of mostly full text, with more and more metadata resulting from Natural Language Processing [19].

Beside the optimization of information retrieval, another research stream which interests automatic story generation focuses on how to present the retrieved data. Since KGs are aligned with ontologies, of which they are the graph realization [20], their "flatness" does not conceal a *flat ontology*, but *conserves* the hierarchical structure of the underlying one. *Design knowledge* determines how the semantics and presentation structure are expressed in the multimedia presentation. In traditional Web environments, this type of design knowledge remains implicit, but Semantic Web technology can be used to model design knowledge explicitly, and to turn annotated media items into structured presentations[ibidem]. A schema which would constitute an ideal starting point for story generation because of the rich logic specification including also temporal tenses is the CIDOC conceptual reference model (CIDOC-CRM) [21]. Currently the effort is being performed of finding strategies, including also ML- approaches, to align already existing knowledge silos with this schema [22].

The Europeana Data Model (EDM)[6], the schema underlying the chosen dataset, is widely compliant with CIDOC-CRM, although the former lacks of some equivalent classes of the latter, whose more fine-grained ontology would better allow to capture interesting nuances for the story building.

2. Related work

Before tackling the literature which is strictly related to the the proposed pipeline, it is desirable to overview contextual information which has played a decisive role in its engineering.

In both education and narration, as well as in Information Retrieval, the *analogy* represents an outstanding device to create interesting associations: the issue of overlapping

[6]Definition of the Europeana Data Model v5.2.8.

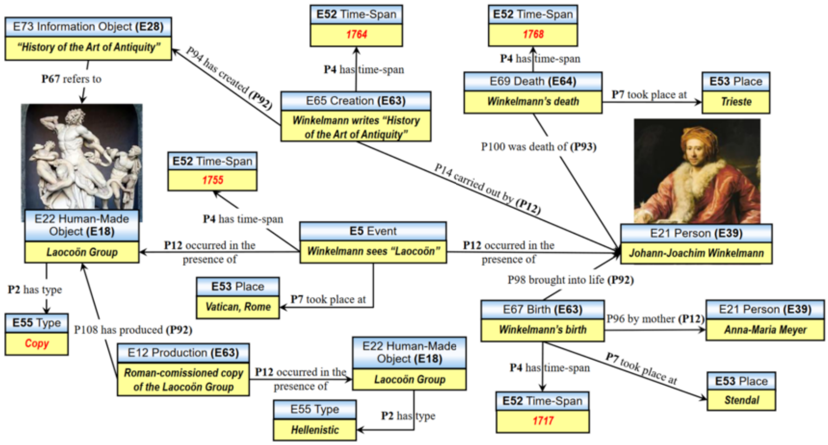

Figure 1. An example of KG based on the CIDOC-CRM schema [21].

concepts throughout different media and the human senses reserved for their fruition, is addressed from a cognitive as well as computational perspective in the *Conceptual Bleinding* theory [23]. According to it, a process starts by finding a partial mapping between elements of two input spaces that are perceived as analogous with respect to their graph representation. Afterwards the so-called generic mental space, encapsulates the conceptual structure shared by the input spaces, generalising and possibly enriching them. This space provides guidance to the next step of the process, where elements from each of the input spaces are selectively projected into a new mental space, called the *blend space*.

Graph representations of concept blends are useful for automated analysis and further processing, but are not very suitable and appealing for human perception of the blended spaces [24]. To improve on this aspect of conceptual blending, algorithms for visual blending and for textual representation of concept graphs have been developed[7].

As a branch of the broader scope of Automatic Language Generation, ASG cannot ignore the relevant results achieved in the former by means of Large Language Models (LLM) such as GPT-3, GLaM, LaMDA, Gopher, PaLM and Megatron-Turing NLG[8]. Story generation remains a challenge because logical coherence among events must be maintained [25].

Fortunately, a handful of works tackling this issue has already been produced, showing that LLMs are a valuable resource which can be easily implemented within a usual pipeline for story generation [26]. Lin & Riedl (2021) propose to evaluate fluency of sequences generated by a blending generation model[9], by using *perplexity* of a base language model, rooted on the intuition that low average perplexity of generated sentences evaluated by the base LLM are consistent with sentences occurring in English, as repre-

[7]The *ConCreTeFlows* workflow collects textual content from two Wikipedia pages about two animals, produces a conceptual map for each of them and creates their blend in three forms: graph-, textual and visual. Application accessible at the website *http://concreteflows.ijs.si/workflow/137/*

[8]https://ai.googleblog.com/2022/04/pathways-language-model-palm-scaling-to.html

[9]It consists of two parts, a language model and a control model, and generates the sentence continuation.

sented by the data used to train the base LLM, which in turn results in seemingly fluent sentences [27]. To address this issue a Story generation with Reader Models (StoRM) [28] is introduced, a framework in which a reader model represented as a knowledge graph infers what a human reader believes about the concepts, entities, and relations regarding the fictional story world (hence, how to progress stories in a plausible way). In Fan & al. (2018) [29] further improvements are gained with a novel form of model fusion that improves the relevance of the story to the prompt, and adding a new gated multi-scale self-attention mechanism to model long-range context.

An alternative approach to this task is constituted by a *reward-shaping technique* that analyses a story corpus and produces intermediate rewards that are back-propagated into a pre-trained LLM in order to guide the model towards a given goal. This method depends on two main models: a LLM based on GPT-2 and a policy model trained via reinforcement learning to select alternative continuations that progress the story incrementally toward the goal [30,31].
New stories can be generated also by reusing existing ones matching a given user query [32]. In Gervas (2005) the plot structure is obtained by a case-based reasoning (CBR) process over a case base of tales and an ontology of explicitly declared relevant knowledge. The resulting story is generated as a sketch of a plot described in natural language by means of NLG techniques.

Extremely relevant topic for my thesis is the *(semi-)automatic world building*, since the initial semantic cloud serves as foundation to the event sequence, and then to the story derivation. Using existing story plots as inspiration, in [33] is described a method that extracts a partial knowledge graph encoding basic information regarding world structure, which is automatically completed utilizing thematic knowledge and used to guide a neural language generation model that fleshes out the rest of the world.
In Yang & Tiddi (2020), the combination of knowledge graphs to provide context clues and implicit knowledge with LLMs is explored, showing that knowledge extracted from KGs can be injected into the stories automatically generated by the LLM [26]. External knowledge graphs are used as well in DICE [34] to provide context clues and implicit knowledge to generate coherent and creative stories.

2.1. Neurosymbolism

Nowadays ASG [b] heavily relies on subsymbolic (e.g. neural or *connectivist*) techniques, which in the best cases succeed to deliver human-like results [35]. On the other hand, symbolic (e.g., rule-based) algorithms come usually into play for macro-planning the events sequence, and in some cases for micro-planning on the lexical or sentence-level, as in [36]. The exclusive use of one or the other approach leads to some pitfalls, as it will be better explored in the dedicated section [2.1].
Neurosymbolic AI, a novel approach concerned with the integration of *reasoning* and *learning*, usually take unstructured data as input, trying to *inject* rules into the neural network by different means, such as, for example, *real logics* [37]. Nevertheless it has been recently thought to be applied to the Semantic Web, for example by using knowledge graphs embeddings on LOD [38], or Deep Learning for Deductive Semantic Web Reasoning [39]; in [40] Semantic Web Knowledge is proposed as a conductor of pretrained LLM. In [41] the best strategy for data integration for Neuro-symbolic NLP is the intersection of the three spaces: *continuous feature vector space, discrete semantic sym-*

bolic space, and *continuous quasi-semantic vector space*[10]. Neurosymbolic automatic story generation (NASG) has already been the central theme of a variety of papers. As in Yao, 2019 [43], where a symbolic plot-control phase is alternate with a neural phase of text generation, my work adds this very methodology to the already mentioned use of semantic enriched data as a way of broadening the interpretation of ”neurosymbolic”, encompassing until now mainly techniques to *inject* rules directly into neural networks. In my thesis, ”neurosymbolic” means properly neural-symbolic, because explicit connections such as semantically enriched links are indeed rules. This awareness may trigger the capability of a system not only of integrating these two approaches, but also to switch between the two at given moments, recognised as pivotal, which would mean briefly assessing for a given task which approach would yield better results. A generative approach for incorporating global structure in the form of relational constraints between different subcomponents of an example can infer the relational constraints present in the training data and then learn a generative model based on the resulting constraint data [44].

On the level of automatic language generation, in Hu & al. (2018) [45] plausible sentences conditioned on representation vectors which are endowed with designated semantic structures are produced using variational auto-encoders. Nye & al. (2021) [46] understand human reasoning as an interplay between two systems: the intuitive/associative and the deliberative/logical one[11]. Neural sequence models exhibit the advantages and failure modes of the first system: they are fast and learn patterns from data, but are often inconsistent and incoherent. Therefore in their work is assessed how candidate generations from a neural sequence model are examined for logical consistency by a symbolic reasoning module, which can either accept or reject the generations. Following results in robust story generation and grounded instruction-show that this approach can increase the coherence and accuracy of neurally-based generations.

The neural-symbolic integration will enable the harnessing of the final output with nuances that only a deterministic model is able to provide, such as the embedding of narratological rules [48], of relations to other similar content, and of relevance according to a given topic. Generated content will thus not only be realistic or generally entertaining (as the current state of the art is already able to deliver) but also fully purposeful, according to whatever is defined by and encoded through the given rules to be the *purpose*.

2.2. From input to graph

In my research I tackle the problem of generating a knowledge graph from keywords, moving from the assumption that the user is unaware of possible interesting connections among some elements, that he is obliged to use as thematic constraint.

The first identified task is therefore KG expansion with *Human in the Loop* (HIL) [49]. In [50] Hyvonen and colleagues present a knowledge-based approach for finding

[10] We forward the reader to [42] for a detailed survey of the main issues, peculiarities and drawbacks occurring when ML methods are adopted in the SW field, in particular regarding semantically enriched embedding models. These models exploit the graph structural information and properties, as well as the additional knowledge available, when rich representation languages as RDFS and OWL are employed. A complementary research direction focused on the preprocessing of LOD for Machine Learning processing exploits vector space embeddings for propositional feature vector representation of RDF data collections [ibid.].

[11] In practice, as a *neurosymbolic* system.

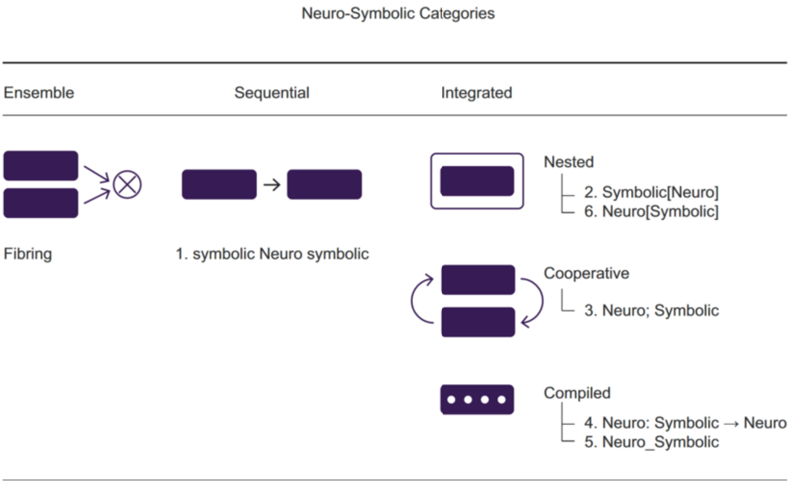

Figure 2. Proposed Neuro-Symbolic Artificial Intelligence categories as in [47].

serendipitous semantic relations between resources in a knowledge graph[12]. The developed system takes two elements as input, and deploys the semantically labelled shortest path between the two.

The RDFsim measure [51], an interactive similarity-based browsing system that exploits knowledge graph embeddings to enable the user to browse the most similar entities of the researched ones, can be intended as a starting point to model *interestingness*, which can be considered as a sort of inverse function for similarity [52].

A semi-automatic workflow to produce story maps from textual documents is assembled in [53], whereby natural language processing and Wikidata services are leveraged to extract key concepts, assemble a logically-ordered sequence of enriched story-map events, producing an interoperable Linked Open Data semantic knowledge base for event exploration and inter-story correlation analyses. This topological interpretation of ”story” matches pretty well the graph-based visualization, which in our case is used as a bridge towards the final full textual realization, but at the same time represents an independent pipeline block, already fully exploitable for educational and creative purposes, a topic deepened in *K12EduKG* [54] as well.

Leveraging heterogeneous domain-specific educational data, K12EduKG extracts concepts and identifies implicit relations with high educational significance. More specifically, it adopts Named Entity Recognition (NER) techniques on educational data like curriculum standards to extract concepts, and employs data mining techniques to identify the cognitive prerequisite relations between educational concepts [ibid].

2.3. *From graph to sequence*

One of the key requirements to facilitate the semantic analysis of historical events in the Web, in the news and in social media is the availability of reference knowledge repositories containing comprehensive representations of events, entities and temporal relations [55]. Existing knowledge graphs, with popular examples including DBpedia, YAGO and

[12]http://www.kulttuurisampo.fi/ff.shtml .

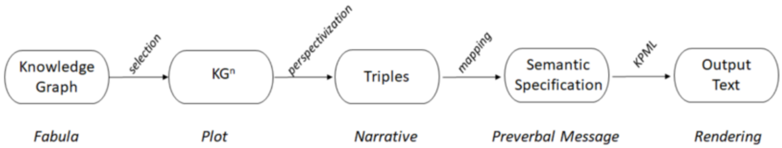

Figure 3. The steps of converting a fragment of a *Fabula* into narrativized natural language text (adapted from [57]).

Wikidata, focus mostly on entity-centric information and are insufficient in terms of their coverage and completeness with respect to events and temporal relations. *EventKG* is a multilingual event-centric temporal knowledge graph that incorporates over six-hundred-ninety thousand events and over two million temporal relations obtained from several large-scale knowledge graphs and semi-structured sources and makes them available through a canonical RDF-representation. Furthermore, narrative overlays together with adequate bindings allow to effectively fuse knowledge and improve retrieval and discovery tasks by structurally aligning underlying repositories driven alone by some narrative [56]. In [57] we see the conversion of a given knowledge graph into natural language as the construction of a narrative about the assertions made by the KG; therefore a pipeline is proposed that can be applied to produce linguistic narratives from knowledge graphs using corresponding rules turning them into a semantic specification for natural language generation [57].

Sequential pattern mining from spatio-temporal data has received much attention in recent years due to its broad application domains such as targeted advertising, location prediction for taxi services, and urban planning. For instance, in [58], an algorithm for mining spatio-temporal event sequences (STESs) from trajectory-based event instances is introduced, which considers each instance to be associated with an event type.

2.4. *From sequence to text*

Most of the previous work on neural text generation from graph-structured data relies on standard *sequence-to-sequence* methods[13]. These approaches linearise the input graph to be fed to a recurrent neural network. Marcheggiani and colleagues propose an alternative encoder based on graph convolutional networks that directly exploits the input structure, showing results on two graph-to-sequence datasets that empirically show the benefits of explicitly encoding the input graph structure [59].
On the same research line, in [60], a neural modelling framework is proposed that jointly learns to generate topically coherent and informative text by computing the representation of the input knowledge graph for each sentential context, and to generate text in a sentence-by-sentence order to improve tractability for long sequence generation.

Another approach considers to first build a document-level path for each output text with each sentence embedding as its node, and a revised self-organising map (SOM) is proposed to cluster similar nodes of a family of document-level paths to construct the directed semantic graph. Then, three subgraph-alignment methods are proposed to extract the maximum matching paths or subgraphs. These directed subgraphs are considered to

[13] https://google.github.io/seq2seq/ .

well preserve extra but relevant content to the short input text, and then they are decoded by the employed pre-trained model to generate coherent long text [61]. Finally, a further suitable solution for interfacing sentence and event level is the use of a Controlled Natural Language (CNL) [62], placing itself between natural and formal languages. This paper proposes the use of CNL for expressing every storytelling system knowledge as a collection of natural language sentences.

3. Research objectives

The Knowledge Graph is a *Directed Acyclic Graph* (DAG) whose nodes and links (respectively, *entities* and *relations*) are semantically enriched, i.e. decorated by textual labels. Building a *narrative* upon a KG, and not a mere textual rendering of the same (i.e., a description) cannot exult from building the initial Knowledge Base with a narrative purpose. Starting from this consideration, and from the knowledge gaps detected in the previous literature review, the following steps are individuated as a necessary and at the same time realistic objective to be tackled in my research project, as well as an homogeneous pipeline proposal:

1. Interest-based Semantic Hub construction;
2. Interestingness-based Knowledge discovery;
3. Event sequences extraction from KG's narrative clots;
4. Link-aware recursive event generation;
5. Metrics and evaluation.

3.1. Interest-based Semantic Hub construction

In the scenario where the user has no precise question, but simply wants to explore links among elements (as it often occurs, for example, in creative writing or narrative generation for art exhibitions) a HIL approach would be necessary, with the goal of constructing a semantic hub which embraces and expands the input concepts. If two elements are given, the most intuitive way to solve this problem is running a *shortest path algorithm* (as the Djikstra's), or finding out the extension contemplating between the nodes in the related KG. On the other hand, if the elements are more, the problem may become very complex. The assumption of this module, is that the vague interests of the user progressively unfolds towards sharper questions, as long as the Knowledge Graph develops in front of him. To allow this dynamic interaction with the user, this module will be implemented as a web application, as well as the other pipeline's modules requiring a similar feature. Given the experimental focus of the set-up, the prototype will allow to select the initial elements only among the default database (i.e. *Europeana*) which will be queried by means of CIPHER on Neo4J[14], connected to the Python[15] framework through a special package. Once the hub embedding the connections among the selected elements is displayed, the process of pruning, rearranging and expanding will proceed. The process will end, when the user is satisfied with the semantic hub he cooperated to generate, starting from which he can independently continue its creative process, or make use of

[14] https://neo4j.com/
[15] https://www.python.org/ .

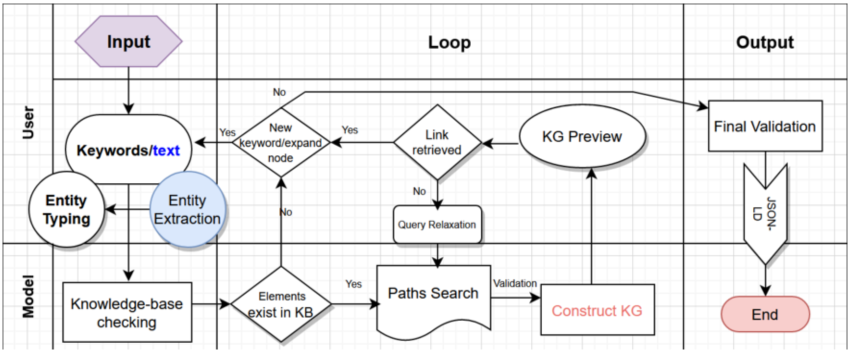

Figure 4. Module 1: HILD&GARD - Interest-based Knowledge Graph (Human In the Loop Data Extraction & Graphically Augmented Relation Discovery).

the next module[16]. This pipeline is intended to expand the knowledge of the user, and although some inspiring unexpected links might already emerge, this is not the main aim of the process.

The proposed pipeline fully rely on already existing techniques. However, to our best knowledge, the formulated task/use case scenario has never been addressed as such.

3.2. *Interestingness-based Knowledge discovery*

The foundation of all upcoming processes, needed to be applied in order to improve the present state-of-art, is the expansion of the initial KG represented by the input constraints and their immediate and obvious semantic connections. The operative field for our research would therefore concern the direction towards which this automatic arborescence shall be directed, applying insights from the conceptual blending theory so that not only plausible, but also interesting connections may be retrieved and visualized.

The rationale of this block is that many of the automatically generated stories until now are impossible to be exploited in practical endeavours. Hence, before proceeding to the machine learning- supported natural language generation module, it is key to establish a tool for the visualization of interesting content, i.e. interesting links in a defined semantic space [50].

The displayed process relies on the previous stage for the retrieval of a consistent semantic hub. A reliable resource for the construction of this block is the concept of ”trivia”, i.e. any fact about an entity which is interesting due to its unusualness, uniqueness or unexpectedness [63]. They would be at first retrieved by means of *web scraping*, then connected into a knowledge graph as in the previous module [3.1]. Alternatively, interestingness will be identified as a sort of mean between similarity and its opposite, so that the displayed divergence is prevented to degenerate in confusion or randomness [51], and the subsequent Knowledge Graph would be further expanded by leveraging ConceptNet5.5 [64] to decorate each node by its semantic cloud.

[16]The related code can be retrieved at https://github.com/Glottocrisio/IKG4CH

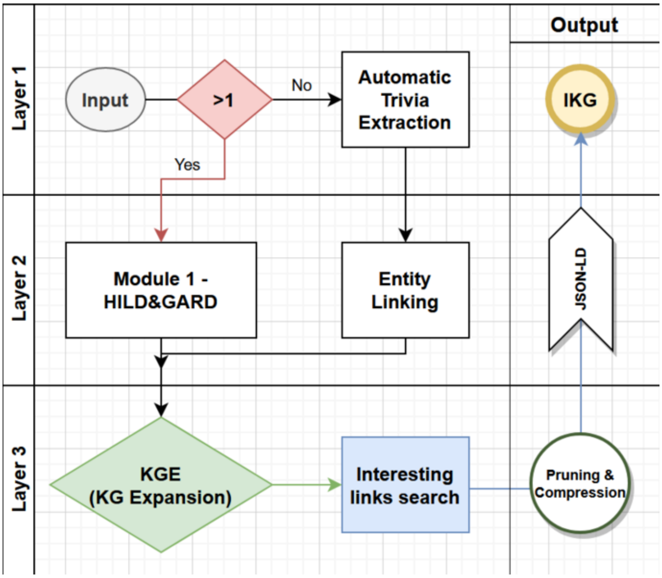

Figure 5. Module 2: IKG - Interestingness-based Knowledge Graph. A pipeline proposal for an Interestingness-based Knowledge Graph.

4. Concluding remarks and future work

The multiplicity of imaginable purposes in the domain of CH is such, that it is unavoidable to wish some sort of granularity in our tool. Whether the story be long or short, be strongly rooted in the input data or have a higher freedom-degree, be thought for memorization, education or enjoyment, the system shall be able to capture it. Therefore, a different model shall be built for each possible combination of these modalities.

Furthermore, the consideration that modern stories are not pure concatenation of events (but also contain dialogues, various story-teller focalizations, digressions, flashbacks, descriptions, and so on) be not yet covered by any system in existence, let us pursue a model which is not only causal, action-based and goal-oriented, but can embed also aesthetic principles and narratological rules. In this respect, the use of neurosymbolism is expected to be used as principal approach to direct text generation.

The following steps to be undertaken immediately after the implementation of the first two, are the ones listed at the beginning of the previous section.

After defining *narrative clots* in the KG, i.e. narratively relevant hubs, a ordered sequence of events shall be extracted from the selected one, which will thereafter constitute the *fabula* in which the story is grounded. The existent related techniques shall be reshaped to be applied to the definition of event, that we have found better suitable, because it better aligns to the current narratological terminology. The DAG is considered to be the *trait-d´union* between the graph and the sequence. The actual state of the art does not take into account nuances such as the synergies among events, and the requirement of recalling previous events, details and characters in an elegant (not redundant) way, which makes this research direction an interesting one to be explored. Beyond the *event-2-event* and the *event-in-event* generation, we recall that also an optimum between interestingness and long-term coherence shall be researched in the frame of this subtopic, which in turn shall be tested on longer outputs (All output examples of treated literature encompass only short-stories).

The task of setting up an evaluation to benchmark the achieved work against the state-of-art results belongs indeed to a later moment of my research. Nonetheless it is necessary to start getting acquainted with patterns occurring in the mathematical formulation of non-strictly mathematical phenomena [52]. In addition to the benchmarks usually deployed to test any model [65], a strong focus will be set on *faithfulness*, *elasticity* and *interestingness*, i.e., respectively, the degree of similarity among different outputs according to the same input, the capability of the system to embrace new input without drastically changing the old output, and the degree of entertainment for human users based on the event level and on universal cognitive assumptions.

5. Acknowledgments

I thank my supervisor, Dr. Maria Pia Di Buono, my co-supervisor and Lab director Prof. Johanna Monti, Dr. Raffaele Manna, Prof. Armando Stellato, Prof. Carlo Strapparava, Prof. Marco Lippi, Prof. Rita Cucchiara and Prof. Alessandro Codello for providing me useful insights to narrow down the project's scope and methodology. I thank the COST Action CA18231 for partial financial support. This research is supported by the Italian National PhD program in Artificial Intelligence.

References

[1] Thoma S, Rettinger A, Both F. Knowledge Fusion via Embeddings from Text, Knowledge Graphs, and Images. arXiv; 2017. ArXiv:1704.06084 [cs, stat]. Available from: http://arxiv.org/abs/1704.06084.

[2] Chen S, Liu B, Fu J, Song R, Jin Q, Lin P, et al.. Neural Storyboard Artist: Visualizing Stories with Coherent Image Sequences. arXiv; 2019. ArXiv:1911.10460 [cs]. Available from: http://arxiv.org/abs/1911.10460.

[3] Gonzalez-Rico D, Fuentes-Pineda G. Contextualize, Show and Tell: A Neural Visual Storyteller. arXiv; 2018. ArXiv:1806.00738 [cs]. Available from: http://arxiv.org/abs/1806.00738.

[4] Benjamin W. Das Passagen-Werk. vol. 5 of Gesammelte Schriften. und Hermann Schweppenhäuser RT, editor. Frankfurt am Main: Suhrkamp-Verlag; 1982.

[5] Collins T, Mulholland P, Bradbury D, Zdrahal Z. Methodology and tools to support storytelling in cultural heritage forums. In: 14th International Workshop on Database and Expert Systems Applications, 2003. Proceedings. IEEE Comput. Soc; 2003. .

[6] Oomen J, Aroyo L. Crowdsourcing in the cultural heritage domain: Opportunities and challenges; 2011. Journal Abbreviation: Agora Pages: 149 Publication Title: Agora.

[7] Shaikh A, Kulkarni SB. Natural Language Processing Applications for Tourism Sector. IOSR Journal of Computer Engineering (IOSR-JCE). 2022;22:27-35.

[8] Hellou M, Lim J, Gasteiger N, Jang M, Ahn HS. Technical Methods for Social Robots in Museum Settings: An Overview of the Literature. International Journal of Social Robotics. 2022 Oct;14(8):1767-86. Available from: https://doi.org/10.1007/s12369-022-00904-y.

[9] Subramanian K. Myth and Mystery of Shrinking Attention Span. International Journal of Trend in Research and Development. 2018;volume 5:1-06.

[10] Negrini M, Di Blas N. In: Digital Storytelling for Cultural Heritage: A Modular, Multi-channel, Multi-scenario Approach. Springer International Publishing; 2014. p. 367-75.

[11] Malegiannaki IA, Daradoumis T, Retalis S. Teaching Cultural Heritage through a Narrative-based Game. Journal on Computing and Cultural Heritage. 2020;13(4):27:1-27:28. Available from: https://doi.org/10.1145/3414833.

[12] Katifori A, Karvouni sM, Kourtis V, Perry S, Roussou M, Ioanidis Y. Applying interactive storytelling in cultural heritage: opportunities, challenges and lessons learned. In: Rouse R., Koenitz H., Haahr M. (eds) Interactive Storytelling. ICIDS 2018. Lecture Notes in Computer Science, vol 11318. Springer, Cham; 2018. p. 603-12.
[13] Paolini P, Di Blas N. In: Storytelling for Cultural Heritage. Springer International Publishing; 2014. p. 33-45.
[14] Dimoulas CA. Cultural Heritage Storytelling, Engagement and Management in the Era of Big Data and the Semantic Web. Sustainability. 2022.
[15] Candela G, Escobar P, Carrasco R, Marco-Such M. A linked open data framework to enhance the discoverability and impact of Culture Heritage. Sage. 2016.
[16] Bales ME, Johnson SB. Graph theoretic modeling of large-scale semantic networks. Elsevier. 2005.
[17] Hao S, Tan B, Tang K, Zhang H, Xing EP, Hu Z. BertNet: Harvesting Knowledge Graphs from Pretrained Language Models. arXiv; 2022. ArXiv:2206.14268 [cs]. Available from: http://arxiv.org/abs/2206.14268.
[18] Hu Y, Chapman A, Wen G, Hall DW. What Can Knowledge Bring to Machine Learning? – A Survey of Low-shot Learning for Structured Data. arXiv; 2021. ArXiv:2106.06410 [cs]. Available from: http://arxiv.org/abs/2106.06410.
[19] de Boer V, Meroño Peñuela A, Ockeloen N. Linked Data for Digital History. In: Romero M, Recio, Ruiz MJC, editors. Historiografía digital: proyectos para almacenar y construir la Historia. Universidad Carlos III de Madrid: Anejos de la Revista de Historiografía;. p. 25.
[20] Geurts J, Bocconi S, van Ossenbruggen J, Hardman L. Towards Ontology-Driven Discourse: From Semantic Graphs to Multimedia Presentations. In: Goos G, Hartmanis J, van Leeuwen J, Fensel D, Sycara K, Mylopoulos J, editors. The Semantic Web - ISWC 2003. vol. 2870. Berlin, Heidelberg: Springer Berlin Heidelberg; 2003. p. 597-612. Series Title: Lecture Notes in Computer Science. Available from: http://link.springer.com/10.1007/978-3-540-39718-2_38.
[21] Definition of the CIDOC Conceptual Reference Model;.
[22] CIDOC-CRM and Machine Learning: A Survey and Future Research;.
[23] Turner M, Fauconnier G. The Way We Think: Conceptual Blending and the Mind's Hidden Complexities. Basic Books; Reprint Edition; 2003.
[24] Znidarsic M, Cardoso A, Gervas P, Martins P, Hervas R, Oliveira Alves A, et al. Computational Creativity Infrastructure for Online Software Composition: A Conceptual Blending Use Case. In: Pachet F, Cardoso A, Corruble V, Ghedini F, editors. Proceedings of the Seventh International Conference on Computational Creativity. France: Sony CSL Paris; 2016. p. 371-9. International Conference on Computational Creativity ; Conference date: 27-06-2016 Through 01-07-2016.
[25] Li S. Using language models in causal story generation; 2020.
[26] Yang X, Tiddi I. Creative Storytelling with Language Models and Knowledge Graphs. In: Proceedings of the CIKM 2020 Workshops. Galway, Ireland; 2020. p. 9.
[27] Lin Z, Riedl M. Plug-and-Blend: A Framework for Controllable Story Generation with Blended Control Codes. In: Proceedings of the Third Workshop on Narrative Understanding. Virtual: Association for Computational Linguistics; 2021. p. 62-71. Available from: https://www.aclweb.org/anthology/2021.nuse-1.7.
[28] Peng X, Xie K, Alabdulkarim A, Kayam H, Dani S, Riedl MO. Guiding Neural Story Generation with Reader Models. undefined. 2021. Available from: https://www.semanticscholar.org/reader/03f079c0d6d5b14a3948e55de0f40677d9634338.
[29] Fan A, Lewis M, Dauphin Y. Hierarchical Neural Story Generation. In: Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers). Association for Computational Linguistics; 2018. .
[30] Tambwekar P, Dhuliawala M, Martin LJ, Mehta A, Harrison B, Riedl MO. Controllable Neural Story Plot Generation via Reinforcement Learning. In: Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence; 2019. p. 5982-8. ArXiv:1809.10736 [cs]. Available from: http://arxiv.org/abs/1809.10736.
[31] Pradyumna Tambwekar, Murtaza Dhuliawala, Animesh Mehta, Lara J Martin, Brent Harrison, Mark O Riedl. Controllable Neural Story Generation via Reinforcement Learning;.
[32] Gervás P, Díaz-Agudo B, Peinado F, Hervás R. Story plot generation based on CBR. Knowledge-Based Systems. 2005 8;18(4-5):235-42.
[33] Ammanabrolu P, Cheung W, Tu D, Broniec W, Riedl MO. Bringing Stories Alive: Generating Interactive

Fiction Worlds. ArXive. 2020:7.
[34] Yang X, Tiddi I. Creative Storytelling with Language Models and Knowledge Graphs;.
[35] Dong C, Li Y, Gong H, Chen M, Li J, Shen Y, et al.. A Survey of Natural Language Generation; 2022. ArXiv:2112.11739 [cs]. Available from: http://arxiv.org/abs/2112.11739.
[36] Elson DK. A Platform for Symbolically Encoding Human Narratives. Association for the Advancement of Artificial Intelligence. 2007:8.
[37] Serafini L, Garcez Ad. Logic Tensor Networks: Deep Learning and Logical Reasoning from Data and Knowledge. arXiv; 2016. ArXiv:1606.04422 [cs]. Available from: http://arxiv.org/abs/1606.04422.
[38] Hitzler P, Bianchi F, Ebrahimi M, Sarker MK. Neural-symbolic integration and the Semantic Web. SW. 2020 Jan;11(1):3-11. Available from: https://www.medra.org/servlet/aliasResolver?alias=iospress&doi=10.3233/SW-190368.
[39] Ebrahimi M, Sarker MK, Bianchi F, Xie N, Doran D, Hitzler P. Reasoning over RDF Knowledge Bases using Deep Learning. arXiv; 2018. ArXiv:1811.04132 [cs, stat]. Available from: http://arxiv.org/abs/1811.04132.
[40] Yu J, Zhang X, Xu Y, Lei X, Guan X, Zhang J, et al. XDAI: A Tuning-Free Framework for Exploiting Pre-Trained Language Models in Knowledge Grounded Dialogue Generation. In: Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining. KDD '22. New York, NY, USA: Association for Computing Machinery; 2022. p. 4422–4432. Available from: https://doi.org/10.1145/3534678.3539135.
[41] Zhang B, Zhu J, Su H. Toward the third generation artificial intelligence. Science China Information Sciences. 2023 Feb;66(2):121101. Available from: https://link.springer.com/10.1007/s11432-021-3449-x.
[42] d'Amato C. Mining the Semantic Web with Machine Learning: Main Issues That Need to Be Known. In: Reasoning Web. Declarative Artificial Intelligence: 17th International Summer School 2021, Leuven, Belgium, September 8–15, 2021, Tutorial Lectures. Berlin, Heidelberg: Springer-Verlag; 2021. p. 76–93. Available from: https://doi.org/10.1007/978-3-030-95481-9_4.
[43] Yao L, Peng N, Weischedel R, Knight K, Zhao D, Yan R. Plan-And-Write: Towards Better Automatic Storytelling. arXiv; 2019. ArXiv:1811.05701 [cs]. Available from: http://arxiv.org/abs/1811.05701.
[44] Young H, Du M, Bastani O. Neurosymbolic Deep Generative Models for Sequence Data with Relational Constraints. ICLR 2021; 2021. Available from: https://openreview.net/forum?id=Y5TgO3J_Glc.
[45] Hu Z, Yang Z, Liang X, Salakhutdinov R, Xing EP. Toward Controlled Generation of Text. arXiv; 2018. ArXiv:1703.00955 [cs, stat]. Available from: http://arxiv.org/abs/1703.00955.
[46] Nye M, Tessler MH, Tenenbaum JB, Lake BM. Improving Coherence and Consistency in Neural Sequence Models with Dual-System, Neuro-Symbolic Reasoning. arXiv; 2021. ArXiv:2107.02794 [cs]. Available from: http://arxiv.org/abs/2107.02794.
[47] Hamilton K, Nayak A, Božić B, Longo L. Is Neuro-Symbolic AI Meeting its Promise in Natural Language Processing? A Structured Review. SW. 2022 Nov:1-42. ArXiv:2202.12205 [cs]. Available from: http://arxiv.org/abs/2202.12205.
[48] Martin LJ. Neurosymbolic automated story generation; 2021.
[49] Manzoor E, Tong J, Vijayaraghavan S, Li R. Expanding Knowledge Graphs with Humans in the Loop. arXiv; 2022. ArXiv:2212.05189 [cs]. Available from: http://arxiv.org/abs/2212.05189.
[50] Hyvonen E, Rantala H. Knowledge-based Relation Discovery in Cultural Heritage Knowledge Graphs;.
[51] Chatzakis M, Mountantonakis M, Tzitzikas Y. RDFsim: Similarity-Based Browsing over DBpedia Using Embeddings. Information. 2021.
[52] Hilderman RJ, Hamilton HJ. Knowledge Discovery and Interestingness Measures: A Survey;.
[53] Bartalesi V, Coro G, Lenzi E, Pagano P, Pratelli N. From unstructured texts to semantic story maps. International Journal of Digital Earth. 2023 Dec;16(1):234-50. Available from: https://www.tandfonline.com/doi/full/10.1080/17538947.2023.2168774.
[54] Chen P, Lu Y, Zheng VW, Chen X, Li X. An automatic knowledge graph construction system for K-12 education. In: Proceedings of the Fifth Annual ACM Conference on Learning at Scale. London United Kingdom: ACM; 2018. p. 1-4. Available from: https://dl.acm.org/doi/10.1145/3231644.3231698.
[55] Gottschalk S, Demidova E. EventKG - the Hub of Event Knowledge on the Web - and Biographical

Timeline Generation. arXiv; 2019. ArXiv:1905.08794 [cs]. Available from: http://arxiv.org/abs/1905.08794.

[56] Kroll H, Nagel D, Balke WT. Modeling Narrative Structures in Logical Overlays on Top of Knowledge Repositories. In: Dobbie G, Frank U, Kappel G, Liddle SW, Mayr HC, editors. Conceptual Modeling. vol. 12400. Cham: Springer International Publishing; 2020. p. 250-60. Series Title: Lecture Notes in Computer Science. Available from: http://link.springer.com/10.1007/978-3-030-62522-1_18.

[57] Porzel R, Pomarlan M, Spillner L, Bateman J, Mildner T, Santagiustina C. Narrativizing Knowledge Graphs. In: Proceedings of the International Workshop on Knowledge Graph Summarization; 2022. .

[58] Aydin B, Angryk R. A Graph-Based Approach to Spatiotemporal Event Sequence Mining; 2016. .

[59] Marcheggiani D, Perez-Beltrachini L. Deep Graph Convolutional Encoders for Structured Data to Text Generation. arXiv; 2018. ArXiv:1810.09995 [cs]. Available from: http://arxiv.org/abs/1810.09995.

[60] Kurisinkel LJ, Chen NF. Graph To Coherent Text: Passage Generation from Knowledge Graphs by Exploiting Edge Representations in Sentential Contexts;.

[61] Wang Z, Zhang X, Du H. Building the Directed Semantic Graph for Coherent Long Text Generation. In: Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing. Online and Punta Cana, Dominican Republic: Association for Computational Linguistics; 2021. p. 2563-72. Available from: https://aclanthology.org/2021.emnlp-main.200.

[62] Concepción E, Mendez G, Gervás P. Mining Knowledge in Storytelling Systems for Narrative Generation. In: Proceedings of the INLG 2016 Workshop on Computational Creativity in Natural Language Generation. Edinburgh, UK: Association for Computational Linguistics; 2016. p. 41-50. Available from: http://aclweb.org/anthology/W16-5507.

[63] Fatma N, K CM, M S. The Unusual Suspects: Deep Learning Based Mining of Interesting Entity Trivia from Knowledge Graphs. In: Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence (AAAI-17);. .

[64] Speer R, Chin J, Havasi C. ConceptNet 5.5: An Open Multilingual Graph of General Knowledge. arXiv; 2018. ArXiv:1612.03975 [cs]. Available from: http://arxiv.org/abs/1612.03975.

[65] Chhun C, Colombo P, Clavel C, Suchanek FM. Of Human Criteria and Automatic Metrics: A Benchmark of the Evaluation of Story Generation. arXiv; 2022. ArXiv:2208.11646 [cs]. Available from: http://arxiv.org/abs/2208.11646.